Miroslaw Malek
University of Texas at Austin, Austin, TX 78712-1084
Phone: (512) 471-5704
E-mail: malek@emx.cc.utexas.edu
Contract Title: System Level Fault Tolerance in Parallel and Distributed Computing Systems
Contract Number: N00014-88-K-0543
Reporting Period: 1 July 1988 - 31 December 1993

FINAL TECHNICAL REPORT

SUMMARY

The major thrust of our effort was focused on the theory and practice of responsive (fault-tolerant, real-time) computing in parallel and distributed processing environments.

New efficient methods of system testing have been developed which shorten a multiprocessor testing time by orders of magnitude and, therefore, can be used at system booting (previous techniques were prohibitively long).

A new design framework for responsive computing was designed and is being implemented for validation. This framework is based on consensus which can be used to provide synchronization, reliable communication, fault diagnosis, checkpointing and even scheduling in multiprocessor environments.

We have formalized and quantified the space-time tradeoff for efficient fault recovery. The system model is a graph, and we were especially successful in analysis of meshes and hypercubes.

We developed a new method called naturally redundant algorithms which allows efficient implementation of application-specific techniques.

We also developed and tested a comprehensive formal model for fault-tolerant parallel algorithm design.

We have made significant contributions to the theory and practice of parallel computer network design, analysis and fault tolerance.

We have also proposed a novel approach to searching by combining existing search techniques in order to maximize performance by finding better solutions in a shorter amount of time.

This report consists of a more detailed description of each accomplishment, followed by a relevant bibliography.

A. Topological Testing, Array Testing and Diagnosis.

We have introduced a new concept, topological testing, and demonstrated several applications in the area of multiprocessor testing. Topological testing uses graph theoretic optimization methods such as the Traveling Salesman Problem, the Chinese Postman

Problem, coloring, path covering and partitioning to minimize the test time. The topological testing techniques can be applied to test a system's behavior and its organization at each level of the system's hierarchy; namely, circuit, logic, register transfer, instruction and processor-memory-switch levels. Specifically, the topological testing approach is demonstrated by developing tests for the multistage interconnection network and the hypercube network. Time optimization for the testing of these networks gives very promising results by taking advantage of inherent parallelism and removing test redundancy. Three orders of magnitude improvement is achieved by applying topological testing techniques to the testing of an existing multistage interconnection network.

We have identified a new solution to contention testing. Using concepts of square of a graph and coloring, we have devised optimal algorithms for contention testing in paths (busses), cycles (rings), trees, meshes and hypercubes. The universality and power of the proposed approach has been proved to be useful, not only to system testing, but also in system integration procedures.

We developed efficient methods for testing packet-switched multistage interconnection networks. In addition to testing the data paths and routing capabilities, we provide tests for detecting faults in the control circuitry, including the conflict resolution capabilities. Using a general model of the switch, we constructed testing sequences for the internal functions of the $f \times f$ switch requiring only $O(f^2 2f)$ tests in the case of round-robin priority and $O(f 2f)$ in the case of fixed priority (f is usually a constant that is less than or equal to 8). We also developed algorithms to test the entire network using, at most, twice the number of tests needed to test a switch, independently of the network size, which results in $O(\log N)$ testing time for an N-processor network. We demonstrated that our method achieves higher coverage and several orders of magnitude reduction in the testing time of complex multiprocessor systems when compared to the previous methods.

We also developed an overlapped segmentation method for testing cellular arrays. The method optimizes the number of tests for pattern sensitivity faults in combinational logic.

References

E. E. Swartzlander, Jr., and M. Malek, "Overlapped Subarray Testing for Wafer Scale Integration," Proceedings of the International Conference on Wafer Scale Integration, 335-343, San Francisco, CA, January 1989.

D. Fussell, M. Malek and S. Rangarajan, "Wafer-Scale Testing/Design for Testability," in the book titled *Wafer Scale Integration*, E. E. Swartzlander, Jr. (ed.), Chapter 9, 413-472, Kluwer, Boston, 1989.

M. Malek, A. Mourad and M. Pandya, "Topological Testing," Proceedings of the 1989 International Test Conference, 103-110, Washington, DC, August 1989.

D. Fussell, S. Rangarajan and M. Malek, "Fault Diagnosis of Linear Processor Arrays," *Defect and Fault Tolerance in VLSI Systems*, I. Koren (ed.), Chapter 6, 149-160, Plenum Press, New York, 1989.

S. Rangarajan, D. Fussell and M. Malek, "Built-in Testing of Integrated Circuit Wafers," IEEE Transactions on Computers, C-39 (2), 195-205, February 1990.

M. Malek and B. Özden, "Optimized Testing of Meshes," Proceedings of the 1990 International Test Conference, 627-637, Washington, DC, September 1990.

A. Mourad, B. Özden and M. Malek, "Comprehensive Testing of Multistage Interconnection Networks," IEEE Transactions on Computers, 935-951, August 1991.

S. Rangarajan, D. Fussell and M. Malek, "Efficient Fault Diagnosis of Switches in Wafer Arrays," Proceedings of the 1992 International Conference on Wafer Scale Integration, 341-351, San Francisco, CA, January 22-24, 1992.

E. E. Swartzlander, Jr., and M. Malek, "Overlapped Subarray Segmentation: an Efficient Test Method for Cellular Arrays," VLSI Design, Special Issue on VLSI Testing, Vol. 1, No. 1, 1-7, 1993.

S. Rangarajan, D. Fussell and M. Malek, "Pipelined Diagnosis of Wafer-Scale Linear Arrays," Journal of Parallel and Distributed Computing, Vol. 20, 212-223, February 1994.

B. Responsive Computer Systems Design Framework.

The simple idea of consensus is to share information among a group of processing elements, preferably in a fault-tolerant manner such that the fault-free members of the PE population are able to consistently agree on and produce correct results despite the actions, malicious or not, of the faulty segment of the population. The importance of the problem stems from its omnipresence. This problem is at the core of protocols handling synchronization, reliable communication, resource allocation, task scheduling, reconfiguration, replicated file systems, sensor reading and other functions.

We have just completed an extensive survey of system-level diagnosis and Byzantine protocols and aim at design and implementation of responsive (fault-tolerant, real-time) consensus for system diagnosis. Our approach is designed to handle large heterogeneous distributed-system environments and is based on system partitioning and protocol hierarchy.

The main result is the real time and fault tolerance analysis of diagnosis protocols used towards implementing a responsive system. We examine the use of redundancy management in both time (reassignment of tasks) and space (masking faults by voting on the result), and the tradeoffs involved. Finally, given a real-time system and a schedule of tasks, we can determine what changes need to be made in order for the system to be responsive.

Since fault tolerance is created by the management of time and space redundancy in software, these techniques are applicable to real-time distributed systems. Also, because this is a software technique, the tradeoff between time criticality and throughput is readily managed in various partitions of the system.

[illegible] is a set of protocols which support various functions such as synchronization, fault diagnosis and load sharing for specific fault models. We completed the design of consensus protocol under a crash fault model. We plan to incorporate other models as well.

References

X. Li and M. Malek, "Analysis of Speedup and Communication/Computation Ratio in Multiprocessor Systems," Proceedings of the Real-Time Systems Symposium, 282-288, Huntsville, AL, December 1988.

A. M. Johnson and M. Malek, "Survey of Software Tools for Evaluating Reliability, Availability and Serviceability," ACM Computing Surveys, 20 (4), 227-269, December 1988; translated and reprinted in Japanese, Kyoritsu Shuppan Co., Ltd., publisher, 1990.

M. Malek, "Responsive Systems, A Challenge for the Nineties," Keynote Address, Proceedings of Euromicro '90, Sixteenth Symposium on Microprocessing and Microprogramming, Amsterdam, The Netherlands, North-Holland, Microprocessing and Microprogramming 30, 9-16, August 29, 1990.

M. Malek, "Responsive Systems: A Marriage Between Real Time and Fault Tolerance," Keynote Address, Proceedings of the Fifth International Conference on Fault-Tolerant Computing Systems, Nürnberg, Germany, Springer-Verlag, Informatik-Fachberichte 283, 1-17, September 25, 1991.

Y. Kakuda, T. Kikuno, M. Malek and H. Saito, "Efficient Checkpointing for Protocol Recovery in Communication Systems," Proceedings of the 1991 International Symposium on Communications, 704-707, Tainan, Taiwan, ROC, December 9-13, 1991.

M. Barborak, M. Malek and A. Dahbura, "The Consensus Problem in Fault-Tolerant Computing," Department of Computer Sciences Technical Report #TR-91-40, The University of Texas at Austin, December 1991.

Y. Kakuda, T. Kikuno, M. Malek and H. Saito, "A Unified Approach to Design of Responsive Protocols," Proceedings of the 1992 IEEE Workshop on Fault-Tolerant Parallel and Distributed Systems, 8-11, Amherst, MA, July 6-7, 1992.

M. Barborak and M. Malek, "UNICON - A UNIversal CONsensus for Responsive Computer Systems," Department of Computer Sciences Technical Report #TR-92-36, The University of Texas at Austin, October 1992.

M. Barborak and M. Malek, "Partitioning for Efficient Consensus," Proceedings of the 26th Hawaii International Conference on System Sciences, Vol. 2, 438-446, Maui, January 5-8, 1993.

M. Barborak, M. Malek and A. Dahbura, "The Consensus Problem in Fault-Tolerant Computing," ACM Computing Surveys, Vol. 25, No. 2, 171-220, June 1993.

M. Malek, "A Consensus-Based Model for Responsive Computing," IEICE Transactions on Information and Systems, Special Issue on Responsive Computer Systems, Vol. E76-D, No. 11, 1319-1324, November 1993.

M. Malek, "A Consensus-Based Framework for Responsive Computer System Design," to appear in Proceedings of the NATO Advanced Study Institute on Real-Time Systems, Springer-Verlag, St. Martin, West Indies, October 5-18, 1992.

C. Space-Time Tradeoffs for Efficient Fault Recovery and Fault Tolerance.

We have formalized and quantified the space-time tradeoff for efficient fault recovery. The mathematical model which proves to be most appropriate is provided by the theory of graphs. We model a computer architecture as a graph G and a job mapped into G as a subgraph H. The dependability qualities of such a system with or without a fault are determined by the resiliency triple defined by three parameters: multiplicity, robustness and configurability. Multiplicity indicates the number of jobs represented by H that can be mapped (executed simultaneously) on a graph (system) G. Robustness represents the number of jobs that can be mapped such that each of their tasks is executed on a different processor. The first parameter allows to measure redundancy in space, while the second one corresponds to redundancy in time. Configurability counts the number of ways a particular job can be mapped onto a system G. We have developed algorithms for evaluation of the above parameters in mesh and hypercube networks. We have also proposed algorithms which optimized the reconfiguration of a faulty job with minimum space/time overhead. Our approach explores the inherent fault tolerance of multiprocessor systems and exploits the topological relationship between the systems architecture and the target applications. The polynomial solutions are provided for paths and trees mapped onto mesh and hypercube networks. In general, the problem is equivalent a subgraph isomorphism and cannot be solved efficiently.

References

M. Malek and K. H. Yau, "The Resiliency Triple in Multiprocessor Systems," Proceedings of 1988 Inter-national Conference on Parallel Processing, 351-358, Chicago, IL, August 1988.

F. Harary and M. Malek, "Quantifying Fault Recovery in Multiprocessor Systems," Mathematical and Computer Modeling, Pergamon Press, Vol. 17, No. 11, 97-104, 1993.

D. Naturally Redundant Algorithms.

We have characterized a class of algorithms suitable for fault-tolerant execution in multiprocessor systems by exploiting the existing embedded redundancy in the problem variables. Because of this unique property, no extra computations need be superimposed to the algorithm in order to provide redundancy for fault recovery, as well as fault detection in some cases. A forward recovery scheme is thus employed with very low time overhead. The method is applied to the implementation of two iterative algorithms: solution of Laplace equations by Jacobi's method and the calculation of the invariant distribution of a Markov chain. Experiments show less than 15% performance degradation for significant problem instances in fault-free situations, and as low as 2.43% in some cases. The extra computation time needed for locating and recovering from a detected fault does not exceed

the time necessary to execute a single iteration. The fault detection procedures provide fault coverage close to 100% for faults which affect the correctness of the computations.

References

L. Laranjeira, M. Malek and R. M. Jenevein, Jr., "On Tolerating Faults in Naturally Redundant Algorithms," Proceedings of the Tenth Symposium on Reliable Distributed Systems, 118-127, Pisa, Italy, September 30 - October 2, 1991.

L. Laranjeira, M. Malek and R. M. Jenevein, Jr., "Experimental Evaluation of Techniques for Fault Tolerance," Department of Computer Sciences Technical Report #TR-92-32, The University of Texas at Austin, July 1992.

L. Laranjeira, M. Malek and R. M. Jenevein, Jr., "Space/Time Overhead Analysis and Experiments with Techniques for Fault Tolerance," *Dependable Computing and Fault-Tolerant Systems*, C. E. Landwehr, et al. (eds.), Vol. 8, No. 3, 303-318, Springer-Verlag, 1993 (presented at the 3rd IFIP International Working Conference on Dependable Computing for Critical Applications, Mondello, Sicily, Italy, September 14-16, 1992).

E. Comprehensive Formal Model for Fault-Tolerant Parallel Algorithms.

We developed a comprehensive formal model for fault-tolerant parallel algorithms and a general methodology for designing reliable applications for multiprocessor environments. The model relies on the formalization of fault-tolerant concepts by means of three nested system predicates and on properties ruling their interrelationships. This rigorous framework facilitates the study of the specific properties that enable an algorithm to tolerate faults. The consequence of that is the outline of systematic design techniques that can be utilized to add fault-tolerant properties to algorithms while preserving their functional characteristics. The proposed model also allows for the quantification of the costs of fault tolerance in terms of space and time redundancy, clarifying the tradeoffs which are inherent to the fault-tolerant design process. The model and design methodology are validated by the uniform application of their principles in the study of several well-known fault-tolerant techniques. The analysis of the cost of fault tolerance in each of these techniques points out that the exploitation of natural redundancy, in applications where this property is present, will lead to the design of fault-tolerant parallel algorithms with very attractive cost/benefit ratio.

References

L. Laranjeira, M. Malek and R. M. Jenevein, Jr., "Nest: A Nested-Predicate Scheme for Fault Tolerance," IEEE Transactions on Computers, Vol. 42, No. 11, 1303-1324, November 1993.

F. Parallel Computer Network Design, Analysis and Fault Tolerance.

We developed a comprehensive technique to evaluate the quality of multistage interconnection networks with respect to their combinatorial power. We analyzed both fault-intolerant and fault-tolerant networks and evaluated them with respect to graceful degradation capability.

We also proposed a new cylindrical banyan multicomputer architecture that still has the best-to-date cost × delay product.

References

I. Gazit and M. Malek, "Fault Tolerance Capabilities in Multistage Network-based Multicomputer Systems," IEEE Transactions on Computers, C-37 (7), 788-798, July 1988.

V. Cherkassky and M. Malek, "Partitioning and Permuting Properties of CC-Banyan Networks," IEEE Transactions on Computers, C-38 (2), 274-278, February 1989.

I. Gazit and M. Malek, "On the Number of Permutations Performable by Extra-Stage Multistage Networks," IEEE Transactions on Computers, C-38 (2), 297-302, February 1989.

V. Cherkassky and M. Malek, "A Measure of Graceful Degradation in Parallel-Computer Systems," IEEE Transactions on Reliability, 38 (1), 76-81, April 1989.

M. Malek and E. Opper, "The Cylindrical Banyan Multicomputer: A Reconfigurable Systolic Architecture," Journal of Parallel Computing, 10, 319-327, 1989.

B. L. Menezes, A. M. Johnson, M. Malek, K. H. Yau and R. M. Jenevein, Jr., "Fault Impact and Fault Tolerance in Multiprocessor Interconnection Networks," Quality and Reliability Engineering International, Vol. 8, 485-500, 1992.

F. Harary and M. Malek, "Variations on a Hypercube: Embellishments of Parallel Architectures, Mathematical and Computer Modeling," Pergamon Press, Vol. 17, No. 11, 105-111, 1993.

G. Search Techniques.

We continue our experiments with a hybrid algorithm where multiple algorithms execute the same problem and exchange information. This promising approach, implemented on shared-memory parallel processors, gave us superlinear speedup. We have achieved an order-of-magnitude speedup on a two-processor system. The only way we can explain it is that, in fact, we have created a new algorithm. We are currently implementing the same approach on a multiprocessor developed at Microelectronics and Computer Technology Corporation in Austin, Texas.

We designed a special-purpose processor for efficient processing of a tabu search for a traveling salesman problem. Our speedup analysis indicates almost three orders of magnitude improvement over state-of-the-art general-purpose processors.

References

M. Malek, M. Guruswamy, M. Pandya and H. Owens, "Serial and Parallel Simulated Annealing and Tabu Search Algorithms for the Traveling Salesman Problem," Annals of Operations Research, 21, 59-84, 1989.

V. Karamcheti and M. Malek, "A TSP Engine for Performing Tabu Search," Proceedings of the 1991 International Conference on Applications-Specific Array Processors, 309-321, Platja d'Aro, Spain, September 2-4, 1991.

M. Malek, M. Guruswamy, H. Owens and M. Pandya, "A Hybrid Algorithm Technique," Department of Computer Sciences Technical Report #TR-89-6, The University of Texas at Austin, March 1989; translated and reprinted in Japanese in the book entitled *Genetic Algorithm and its Applications*, Kido, Takashi, (Hiroaki Kitano, ed.), Sangyo Tosho Publishing Co., Ltd., May 1993.

H. Surveys and Reviews on Parallel Computing.

References

M. Malek, "Fault-Tolerant Computing in Europe - An Update," European Science Notes Information Bulletin, ESNIB 91-01, 36-39, 1991.

M. Malek, "Transputer Applications '90," European Science Notes Information Bulletin, ESNIB 91-01, 39-43, 1991.

M. Malek, "Closing a Gap between Parallel Computing and Physical Sciences," European Science Notes Information Bulletin, ESNIB 91-03, 10-13, 1991.

M. Malek, "Starlight - High Performance, Low Cost, Small Package," European Science Notes Information Bulletin, ESNIB 91-04, 65-67, 1991.

M. Malek, "Parallel and Distributed Computing," in the *Encyclopedia of Computer Science and Technology*, A. Kent and J. G. Williams (eds.), Vol. 24, Supplement 9, 295-315, Marcel Dekker, New York, 1991.

M. Malek, "High-Performance Computing in Europe," European Science Notes Information Bulletin, ESNIB 92-08, 471-517, 1992.

M. Malek, "Parallel and Distributed Computing," in the *Encyclopedia of Microcomputers*, A. Kent and J. G. Williams (eds.), Vol. 13, 107-127, Marcel Dekker, New York, 1993.

Invited presentations:

"State of the Art in Parallel Computing," University of Utah, Salt Lake City, Utah, April 14, 1989.

"Topological Testing," University of Minnesota, Minneapolis, Minnesota, April 20, 1989.

"Topological Testing," Technion University, Haifa, Israel, June 2, 1989.

"Fault-tolerant Computing," Samsung Labs, Seoul, Korea, August 1989.

"Topological Testing," ETRI, Seoul, Korea, August 1989.

"Hybrid Optimization Techniques," Seoul National University, Seoul, Korea, August 1989.

"Topological Testing," Kokusai Denshin Denwa Company, Tokyo, Japan, August 1989.

"Hybrid Optimization Techniques," Kokusai Denshin Denwa Company, Tokyo, Japan, August 1989.

"Fault-tolerant Algorithms," Kokusai Denshin Denwa Company, Tokyo, Japan, January 22, 1990.

"Topological Testing," RICOH Company, Ltd., Tokyo, Japan, January 23, 1990.

"Topological Testing," SONY Corporation, Tokyo, Japan, January 23, 1990.

"Performability and Responsiveness of Computer Systems," Japanese Technology Form, Tsukuba, Japan, January 25, 1990.

"Topological Testing," New Mexico State University, Las Cruces, New Mexico, April 6, 1990.

"Responsive Systems, A Challenge for the Nineties," Keynote Speech, Euromicro 16th Symposium on Microprocessing and Microprogramming, Amsterdam, The Netherlands, August 29, 1990.

"Fault-tolerant Algorithms," International Federation for Information Processing 17th Meeting, Les Trois Ilets, Martinique, February 2, 1990.

"A Hybrid Algorithm Technique," The Institute of Management Sciences/Operations Research Society of America 29th Joint Annual Meeting, Las Vegas, Nevada, May 7, 1990.

"A Hybrid Algorithm Technique," Gesellschaft für Mathematik und Datenverarbeitung (GMD), Schloss Birlinghoven, Bonn, Germany, October 15, 1990.

"Problems, Architectures and Implementations," Workshop on Parallel Computing for Physical Sciences, Institut National de Recherche en Informatique et en Automatique (INRIA), Rocquenfort, near Paris, France, December 3, 1990.

"A Consensus Problem, The Key to Fault Management," Management Policy Workshop, Imperial College, The University of London, England, February 27, 1991.

"Can We Design a General-Purpose Parallel Computer?," University of Catalunya, Barcelona, Spain, April 15, 1991.

"Can We Design A General-Purpose Parallel Computer?," European Computer Research Center, Münich, Germany, April 24, 1991.

"High-Performance Computing in Europe," Office of Naval Research European Office, London, England, September 1991.

"Responsive Systems: A Marriage between Real Time and Fault Tolerance," Fifth International Conference on Fault-Tolerant Computing Systems, Keynote Address, Nurnberg, Germany, September 25-27, 1991.

"High-Performance Computing in Europe," Office of Naval Research, Washington, D.C., January 1992.

"Research Funding in the United States," Office of Naval Research European Office, London, England, March 1992.

"Responsive Consensus in Distributed Systems," Keio University, Yokohama, Japan, March 1992.

"Responsive Systems: A Marriage between Real Time and Fault Tolerance," Keio University, Yokohama, Japan, March 1992.

"Research Funding in the United States," University of Warsaw, Poland, March 30, 1992.

"Responsive Systems: A Marriage between Real Time and Fault Tolerance," University of Warsaw, Poland, March 30, 1992.

"Research Funding in the United States," Technical University of Budapest, Hungary, April 1, 1992.

"Responsive Systems: A Marriage between Real Time and Fault Tolerance," Technical University of Budapest, Hungary, April 1, 1992.

"Research Funding in the United States," Czech Technical University, Prague, Czechoslovakia, April 2, 1992.

"Responsive Systems: A Marriage between Real Time and Fault Tolerance," Czech Technical University, Prague, Czechoslovakia, April 2, 1992.

"A Consensus-Based Framework for Responsive Computer Systems," Hiroshima Prefectural University, Japan, April 23, 1992.

"Responsive Computer Systems," Hiroshima University, Japan, April 24, 1992.

"A Consensus-Based Framework for Responsive Computer Systems," Osaka University, Japan, April 28, 1992.

"A Hybrid Algorithm Technique for Improved Performance and Dependability," University of Kyoto, Japan, June 19, 1992.

"A Consensus-Based Framework for Responsive Computer Systems," NTT Software Laboratory, Tokyo, Japan, July 24, 1992.

"The 21st Century Computing - A Quest for Hybridization," Parallel Computing seminar sponsored by the local IEEE Student Computer Society, November 3, 1992.

"A Consensus-Based Framework for Responsive Computer System Design," Texas A&M University, Department of Computer Science, College Station, TX, March 4, 1993.

"A Consensus-Based Framework for Responsive Computer System Design," Old Dominion University, Norfolk, VA, April 16, 1993

"A Consensus-Based Framework for Responsive Computer System Design," GMD FIRST, Berlin, Germany, July 2, 1993.

"A Consensus-Based Framework for Responsive Computer System Design," Max Planck Institute, Berlin, Germany, July 13, 1993.

Contributed presentations:

"Networks Testability and Fault Tolerance," Creating Reliable Products Conference, Palo Alto, California, October 10-14, 1988.

"Analysis of Speedup and Communication/Computation Ratio in Multiprocessor Systems," Ninth Real-Time Systems Symposium, Huntsville, Alabama, December 8, 1988.

"Design and Test of Parallel Computer Networks," NATO Workshop on Networks, Copenhagen, Denmark, June 21, 1989.

Presented a review of research on testing and fault tolerance, Annual Office of Naval Research Review, Washington, D.C., June 29, 1989.

"Topological Testing," International Test Conference, Washington, D.C., August 1989.

"Fault-tolerant Algorithms," International Federation for Information Processing 17th Meeting, Les Trois Ilets, Martinique, February 2, 1990.

"A Hybrid Algorithm Technique," The Institute of Management Sciences/Operations Research Society of America 29th Joint Annual Meeting, Las Vegas, Nevada, May 7, 1990.

"The 21st Century Computing - A Quest for Hybridization," The University of Texas at Austin, Department of Electrical and Computer Engineering, January 28, 1991.

"Can We Design A General-Purpose Parallel Computer?," IEE Parallel Computing Seminar, Special Session on Current Issues and New Results, April 18, 1991.

"A TSP Engine for Performing Tabu Search," 1991 International Conference on Applications-Specific Array Processors, Platja d'Aro, Spain, September 3, 1991.

"Responsive Multicomputer Systems," tutorial given at the 12th International Conference on Distributed Computing Systems, Yokohama, Japan, June 9-12, 1992.

"Practical Open Problems for Graph Theorists in Parallel Computer Networks and Fault Tolerance," Graph Theory Workshop, Keio University, Yokohama, Japan, July 23, 1992.

"Space/Time Overhead Analysis and Experiments with Techniques for Fault Tolerance," The Third IFIP International Working Conference on Dependable Computing for Critical Applications, Mondello, Sicily, Italy, September 14-16, 1992.

"A Consensus-Based Framework for Responsive Computer system Design," The NATO Advanced Study Institute on Real-Time Systems, St. Martin, West Indies, October 5-18, 1992.

"Partitioning for Efficient Consensus," The Twenty-Sixth Hawaii International Conference on System Sciences, Maui, January 5-8, 1993.

"Scalable Consensus and its Reliability Limits," The Twenty-Third IFIP Meeting on Dependable Computing and Fault Tolerance, Islamorada, Florida, January 27-31, 1993.

Honors received:

Texas Atomic Energy Research Foundation Centennial Fellowship in Electrical Engineering, 1987-1989.

Bettie Margaret Smith Professorship in Engineering, The University of Texas at Austin, 1989-1991.

Southwestern Bell Foundation Endowed Professorship in Engineering, 1991-present.

Program Committee Member, Euromicro Workshop on Real-time Systems, Como, Italy, June 1989.

Program Committee Member, Workshop on Wafer-scale Integration, Como, Italy, June 1989.

Program Committee Member, 10th International Conference on Distributed Computing Systems, Paris, France, May 28-June 1, 1990.

Program Committee Member, International Conference on Application-Specific Array Processors, Princeton, New Jersey, September 5-7, 1990.

Workshop on Parallel Computing for Physical Sciences, Institut National de Recherche en Informatique et en Automatique (INRIA), Program Co-Chairman, Rocquenfort, near Paris, France, December 3-5, 1990.

The 21st International Symposium on Fault-Tolerant Computing, Program Committee Member, Montreal, Canada, June 25-27, 1991.

International Conference on Application-Specific Array Processors, Program Committee Member, Platja d'Oro, Spain, September 2-4, 1991.

Parallel Computing '91, Program Committee Member, London, September 3-6, 1991.

Fifth International Conference on Fault-Tolerant Computing Systems, Keynote Speaker, Nurnberg, Germany, September 25-27, 1991.

Tenth Symposium on Reliable Distributed Systems, Program Committee Member, Pisa, Italy, September 30, October 1- 2, 1991.

The Third International Workshop on Responsive Computer Systems, Lincoln, NH, General Chair, September 29 - October 1, 1993.

The Twenty-Fourth Annual International Symposium on Fault-Tolerant Computing, General Chair, Austin, Texas, June 15-17, 1994.

Associate Editor, Journal of Real-time Systems, continuing.

Subject Area Editor, Journal of Parallel and Distributed Computing, continuing.

Prizes or awards received:

ONR Award for Outstanding Performance, June 1991.

Promotions obtained:

Texas Atomic Energy Research Foundation Centennial Fellowship in Electrical Engineering, commencement on September 1, 1987.

Bettie Margaret Smith Professorship in Engineering, The University of Texas at Austin, commencement on September 1, 1989.

Southwestern Bell Foundation Endowed Professorship in Engineering, The University of Texas at Austin, commencement on September 1, 1991.

Liaison Scientist, Office of Naval Research European Office, London, England, commencement on June 15, 1990, with ending date of February 29, 1992.

Graduate students supported:

Michael A. Barborak (Ph.D.)
Guillaume P. Brat (Ph.D.)
Luiz A. Laranjeira (Ph.D.)
Banu Özden (Ph.D.)
Mihir A. Pandya (Ph.D.)
Leslie G. Poer (Ph.D.)
Kitty Hiu Yau (Ph.D.)
Myungchul Yoon (Ph.D.)

Minorities supported:

Luiz A. Laranjeira (Hispanic)
Banu Özden (Middle Eastern)
Mihir A. Pandya (Indian)
Kitty Hiu Yau (Asian)
Myungchul Yoon (Asian)

[redacted]
University of Texas at Austin, Austin, TX 78712-1084
Phone: (512) 471-5704
E-mail: malek@emx.cc.utexas.edu
Contract Title: System Level Fault Tolerance in Parallel and Distributed Computing Systems
Contract Number: N00014-88-K-0543
Reporting Period: 1 July 1988 - 31 December 1993

TRANSITIONS AND DoD INTERACTIONS

A robust implementation of the two-dimensional Fast-Fourier Transform algorithm was mapped onto IBM's GF-11 parallel computer, in addition to topological testing, at IBM's T. J. Watson Research Center in New York.

The Hybrid Algorithm Technique was implemented on Motorola's parallel computer, Pleiades, at Microelectronics and Computer Technology Corporation in Austin, Texas. The new technique offers super-linear speedup for search methods, such as tabu and simulated annealing.

IBM's GF-11 supercomputer controls the interconnection network centrally. This control scheme makes the run-time network configuration very expensive. Graph theoretical approaches were investigated by Banu Özden during the summer of 1990 at the T. J. Watson Research Center in Yorktown Heights, New York, to emulate distributed control on the Benes network of GF-11 parallel computer which reduced the cost of dynamic interprocessor communication.

The Hybrid Algorithm Technique was successfully implemented for the Traveling Salesman Problem on a network of UNIX machines by Mihir A. Pandya during the summer of 1990 at IBM, Austin, Texas. The distributed processing implementation used Remote Procedure Calls and incorporated real-time and fault-tolerant features so that the execution completed in the designated time, and results were available as long as a single processor was available. Near-linear speedup was obtained on 2-4 processors.

The topological testing approach was used on the Benes network of the IBM's Gigaflop-11 (GF-11) multiprocessor. My student spent the summer of 1989 at the T. J. Watson Research Center in Yorktown Heights, New York, and applied a combination of level-sensitive scan design (LSSD) and topological testing to the GF-11 supercomputer which resulted in significantly higher fault coverage and improvement in test time.

Miroslaw Malek
University of Texas at Austin, Austin, TX 78712-1084
Phone: (512) 471-5704
E-mail: malek@emx.cc.utexas.edu
Contract Title: System Level Fault Tolerance in Parallel and Distributed Computing Systems
Contract Number: N00014-88-K-0543
Reporting Period: 1 July 1988 - 31 December 1993

SOFTWARE AND HARDWARE PROTOTYPES

(None to report)